



Institute for Scientific Computing Research



ISCR Sabbatical Program

# ISCR Sabbatical Program

The ISCR hosts several sabbatical visitors each year. These faculty stay for 3 to 12 months and work on research in close collaboration with other LLNL staff members. Following are reports on sabbatical activities during FY 2003.

Visitor	University	LLNL Contact	Time at LLNL	Page
Constantine Bacuta .....	Penn State University .....	Rob Falgout .....	July 2, 2003 – Sept. 26, 2003 .....	151
Zhiqiang Cai .....	Purdue University .....	Charles Tong .....	May 21, 2003 – Aug. 2, 2003 .....	152
Hans De Sterck .....	University of Colorado, Boulder .....	Robert Falgout .....	Aug. 13, 2003 – Sept. 5, 2003 .....	153
Rod Fatoohi .....	San Jose State University .....	Jeffery Vetter .....	June 16, 2003 – Aug. 22, 2003 ....	154
Leszek Marcinkowski .....	Warsaw University .....	Carol Woodward .....	June 11, 2003 – July 11, 2003 .....	155
Anne Ngu .....	Southwest Texas State University ....	Terence Critchlow .....	June 9, 2003 – Sept. 15, 2003 .....	156
John Trangenstein .....	Duke University .....	Richard Hornung .....	Jan. 3, 2003 – May 2, 2003 .....	157
Ludmil Zikatanov .....	Penn State University .....	Rob Falgout .....	July 2, 2003 – Aug. 29, 2003 .....	159

# Distributed Relaxation Methods for Stokes Problems

**Constantin Bacuta**

Pennsylvania State University

*LLNL Contact*

**Rob Falgout**

## Summary

---

In the context of solving indefinite systems arising from Stokes or Navier–Stokes equations using multigrid methods, a real challenge for proving convergence results is the construction of an appropriate smoother. The distributed relaxation method for the incompressible Stokes problem is based on a change of variables that leads to a lower triangular system with Laplace operators on the main diagonal, for which multigrid methods are more suitable. We propose a finite element formulation of Achi Brandt's distributed relaxation method. Using this approach, we intend to construct an effective and robust multigrid method for solving Stokes-type systems and other elliptic PDE systems, which can be reformulated as saddle-point problems.

In our finite element discretization of the Stokes problem, we introduce a discrete change of variables and study the properties of the transformed problem. We show that under reasonable regularity assumptions, the transformed problem admits almost block lower-triangular form. Here, “almost” represents the fact that the variables are coupled only through boundary values and the truncation error. We pay special attention to the boundary conditions of the new discrete variables. To impose the coupling Dirichlet boundary conditions, we use a stable quasi-interpolant that brings the derivatives of the pressure variable into the space of the velocity variables. We also prove the coercivity of the operator on the main diagonal of the transformed problem. We are in the process of implementing an extensive set of numerical experiments to demonstrate the effectiveness of the new algorithm we propose.

In the future, we will consider the almost incompressible elasticity problem. It can be reformulated as a saddle-point problem, and hence we can use a similar change of variables to end up with an almost lower triangular matrix suitable for smoothing. Another direction could be solving Maxwell equations, also reformulated as a saddle-point problem, and finding appropriate distributive relaxation change of variables and appropriate discrete spaces.

# Computation of Incompressible Newtonian Flows

**Zhiqiang Cai**

Purdue University

*LLNL Contact*

**Charles Tong**

## Summary

---

The primitive physical equations for incompressible Newtonian fluid flows are the conservation of momentum and the constitutive law. The constitutive law relates the stress tensor to the deformation rate tensor and pressure, and it states the incompressibility condition. By differentiating and eliminating the stress, one obtains the well-known second-order incompressible Navier–Stokes equations in the velocity–pressure formulation.

A tremendous amount of computational research has been done on these equations, but they may still be one of the most challenging problems in computational fluid mechanics. This summer, we studied two issues that are concerns of scientists at Lawrence Livermore National Laboratory: coupling between the momentum equation and incompressibility condition and unknown outflow boundary conditions in some applications.

The velocity–pressure formulation of incompressible Navier–Stokes equations is a coupled second-order partial differential system. The coupling dramatically complicates numerical procedure. First, spatial discretization requires a stable pair of finite element spaces to approximate the velocity and pressure. (Neither the Q1Q0 and Q1Q1 trial spaces are stable and have to be artificially stabilized.) Second, saddle-point problems at each time step due to implicit or semi-implicit time integration are very expensive to solve. To overcome such a difficulty, a popular approach dating back to the 1960s is to use splitting (projection) methods. The splitting method decouples a system through approximation in time, and hence the solution loses accuracy. More specifically, pressure boundary conditions used in splitting methods are artificial and inconsistent, causing inaccurate approximation and giving rise to a numerical boundary layer. We fixed this problem by a modification of the pressure calculation. An alternative is to use numerical methods, which we are developing, based on the stress–velocity formulation. Time integration of this formulation leads to a triangular system, which is decoupled. So we can calculate first the stress and then the velocity in a very efficient way.

The incompressible Navier–Stoke equation is an initial-boundary value problem in mathematical terms. This implies that one needs boundary conditions everywhere for its wellposedness. Unfortunately, many applications can be modeled by incompressible Navier–Stoke equations that have outflow boundary conditions that are unknown. This is a computational modeling issue rather than a purely numerical issue. Without physical intuition, one could use information from the differential equations themselves to approximate outflow boundary conditions. Nevertheless, it is impossible to get correct boundary conditions when using any traditional numerical approach. To compensate for incorrect outflow boundary conditions, one has to use a large computational domain. To tackle this problem, we are studying least-squares methods based on first-order partial differential equations because the least-squares formulation automatically employs the original first-order differential system as outflow boundary conditions, which are consistent. Previous computational experiments on the backward-facing step flow in a long channel show that calculation based on a least-squares method can be done accurately on a very small computational domain.

# Parallel Algebraic Multigrid Solvers—Least Squares (AMS—LS) Applied to Hyperbolic Partial Differential Equations (PDEs)

**Hans De Sterck**

University of Colorado, Boulder

*LLNL Contact*

**Robert Falgout**

## Summary

The objective of this project is to apply parallel algebraic multigrid (AMG) solvers to least-squares (LS) finite element (FE) discretizations of partial differential equations (PDEs) to obtain a solution method for these PDE problems that scales well on large parallel computers. In particular, we investigate the scalability of the AMG—LS approach for hyperbolic PDEs. Hyperbolic PDEs arise in many domains of science and engineering, including fluid dynamics and plasma dynamics, with applications including aerospace engineering, nuclear fusion, and astrophysics. In many of these applications, the problems in need of solution are extremely large. Large parallel computers are available, but present-day methods are not scalable to them.

The first part of the project was to couple the serial LS FE code FOSpack, which was developed at the University of Colorado (CU) at Boulder by John Ruge, to the parallel solver Library Hypre, which was developed at Lawrence Livermore National Laboratory's Center for Advanced Scientific Computing in a project directed by Rob Falgout. FOSpack, which is written in Fortran, has been coupled with the Hypre library, using Hypre's Fortran calling interface. The purpose of this coupling is threefold. First, the coupling provides an infrastructure with which to compare the FOSpack and Hypre AMG solvers. Second, it makes the other Hypre solvers available for use in FOSpack. Third, the coupling makes all the FOSpack features (for example, advanced equation parser, adaptive grid refinement, full MG, nonlinearity, higher-order elements) available for tests with the Hypre solvers. For the second part of the project, a preliminary parallel driver for Hypre has been written that implements the LS discretization of linear hyperbolic PDEs. It had been found before that an augmented equation system scales well with a serial AMG code, using a slightly modified AMG interpolation formula that takes into account the sign of the off-diagonal matrix elements. We have investigated parallel scaling of this formulation using the BoomerAMG solver provided in Hypre. The modified interpolation was implemented in BoomerAMG by Ulrike M. Yang.

It is the goal of the CU Boulder scalable solvers group to make the FOSpack code parallel. Coupled with Hypre, this will allow scientists and engineers to research parallel scalable solvers for a wide variety of LS-discretized PDE problems, including elasticity, fluid dynamics, and coupled systems, by using features such as adaptive grid refinement and higher-order elements. A parallel FOSpack—Hypre code will allow solving large complex science and engineering problems in a scalable way. For the particular application of hyperbolic PDEs, further large-scale parallel scaling tests will be performed using Hypre—BoomerAMG, and scalable AMG algorithms for hyperbolic PDEs will be further developed.

# Performance Evaluation of High-Speed Interconnects Using Microbenchmarks and Scientific Applications

**Rod Fatoohi**

San Jose State University

*LLNL Contact*

Jeffrey Vetter

## *Summary*

---

Our motivation is to study the affect of high-speed interconnects on the performance and scalability of scientific applications.

We are developing a set of microbenchmarks to characterize different aspects of the interconnects. Besides the traditional network parameters (latency and bandwidth), we employ parameters influenced by the network topology (such as bisection bandwidth and network diameter). Our microbenchmarks are implemented using two network APIs: the message passing interface (MPI) and the BSD sockets interface. We also employ several MPI-based scientific applications that use a variety of communication patterns. On this project, we are using the ASCI Purple benchmark codes.

Our microbenchmarks and applications have been tested on several systems using different interconnects, including two IBM SP switch-based systems (Frost and Blue) and two Quadrics-based clusters (pengra/MCR and TC2K). Future testing includes Cray XL, an InfiniBand-based cluster, and BlueGene/L.

# An Overlapping Schwarz Method for Nonlinear Unstructured Finite Element Elliptic Problems

**Leszek Marcinkowski**

Warsaw University

&

**Xiao-Chuan Cai**

University of Colorado, Boulder

*LLNL Contacts*

**Panayot S. Vassilevski**

&

**Carol S. Woodward**

CASC

## Summary

We are interested in solving the nonlinear algebraic equations typically arising in the situation of nonlinear second-order elliptic PDEs discretized using finite elements. These problems are very difficult to solve as the nonlinearities involved are often unbalanced.

Few years ago Cai and Keyes proposed a new approach for solving nonlinear problems called Additive Schwarz Inexact Newton Method (ASPIN), which replaces the original nonlinear system of equations with a new one which generally has less difficult nonlinearities. ASPIN generates nonlinear subproblems that can be solved in parallel. The outputs of these subproblems are used to construct a global nonlinear system, which is solved by an outer Newton method. Therefore, Newton's method on small sparse problems in the original coordinates is used inside of a primary Newton loop on a transformed problem that possesses a (generally) dense Jacobian. The triumph is that the linear problems generated by the outer Newton method on the transformed problem can be preconditioned by sparse solves on subdomains of the original problem.

However for finely resolved problems on a large number of processors there is a need to introduce a coarse grid which would possibly reduce the number of linear iterations on each Newton step. This is challenging since there is no clearly best way to construct a coarse problem for the transformed system. Therefore we wanted to combine this approach with the concept of subspaces that are defined over the whole domain. The mesh is fine in every local subdomain and is gradually coarsened away from it. Such spaces were proposed first for linear problems by R.E. Bank and M. Holst. In the new setting we expected that we would not have to use additional coarse grid to reduce the number of linear iterations in Newton method.

The objective of the present paper is to construct a coarser version of the original fine grid nonlinear problem based on algebraic means targeting nonlinear problems discretized on generally unstructured grids. The coarsening is carried out outside a given domain  $G$  which is assumed to be a union of fine grid elements. The method is a generalization of the element agglomeration AMGe.

We confirmed that our algorithm is convergent with a rate bounded independently of the mesh size, i.e., the number of nonlinear Newton iterations and linear iterations remain constant independently of the mesh and they also remain fairly insensitive with respect to the number of subdomains. For our experiments we used 'hypre' package and the graph partitioner software 'METIS'. We also made progress in developing theory which will demonstrate that for a class of nonlinear problems our Newton method is locally convergent.

# Automatic Discovery and Interaction with Bioinformatics Web

**Anne H.H. Ngu,**  
Southwest Texas State University  
&

**Daniel Rocco**  
Georgia Institute of Technology  
*LLNL Contact*

**Terence Critchlow**

## Summary

---

The World Wide Web provides an incredible resource to genomics researchers in the form of query access to distributed data sources—e.g., BLAST sequence homology search interfaces. The number of these autonomous sources and their rate of change outpace the speed at which they can be manually classified, meaning that the available data is not being used to its full potential. Manually maintaining a wrapper library will not scale to accommodate the growth of genomics data sources on the Web, challenging us to produce an automated system that can find, classify, interact, and wrap new sources without tedious and low-level coding of wrappers. Previous research has not addressed the problem of automatically locating, classifying, and integrating classes of Bioinformatics Web data sources. A correct classification of any kind of Web data source must address both the domain of the source and the conversational/interaction semantics that is inherent in the design of many of the Web sources. In this paper, we proposed a novel approach to classification of Web data sources that takes into account both the capability and the conversational semantics of the source. The goal of my summer assignment is to enhance the current classifier at LLNL with the ability to discover the interaction pattern of a Web source that will lead to increased accuracy of the classification process. At the same time, it enables the extraction of process semantics that are necessary for the automatic generation of wrappers that can interact correctly with the sources.

I first identify the characteristics of the interaction pattern for BLAST data sources. Two algorithms are implemented for identifying the common interaction patterns. Three shortcomings are identified in the first heuristics-based algorithm. A second more robust and efficient algorithm based on computing page difference (PageDiff) between two HTML pages is proposed. A set of experiments is conducted that demonstrate that the PageDiff approach outperforms heuristics by a factor of four. Integrating the PageDiff algorithm with the current LLNL classifier resulted in about a 10% increase in the number of correctly classified sources.

In the future, I want to pursue two goals. The first one is the identification of more complex interaction patterns that require human input. The second is the optimization of BLAST queries. Given that many BLAST sources could answer a particular query, which source should be selected at runtime? How is correlation between sites found out during the classification process? Can the quality of service be inferred during the classification process and used for finding out the correlation? If funding is available, I plan to follow up the collaboration with future visits to Livermore. I also plan to apply for a joint grant with the Data Science Group in the area of efficient and reliable access to Web sources.

# Adaptive Mesh Refinement for Viscous Fingering in Porous Media

**John Trangenstein**

Duke University

*LLNL Contact*

**Richard Hornung**

## Summary

During his sabbatical visit to LLNL, John Trangenstein worked on porting his adaptive mesh refinement code to several LLNL clusters, on extending the modeling, and on improving the algorithms used in this code.

Together with Andy Wissink and David Hysom of the Center for Applied Scientific Computing (CASC), Trangenstein adapted the Structured Adaptive Mesh Refinement Application Infrastructure (SAMRAI) BinaryTree, BoxTopology, and SpatialKey C++ classes to his code. These algorithms are useful in reducing the computational complexity of regridding, determining the communication between grid patches, and assigning patches to processors in order to reduce the need for communication.

Conversations with Robert Falgout and Panayot Vassilevski were useful in both debugging and improving the performance of multigrid algorithm in Trangenstein's porous media Adaptive Mesh Refinement (AMR) code. By incurring some additional communication costs in the multigrid smoother, he was able to reduce the number of multigrid iterations by about 25%. Trangenstein also completed the debugging of the 3D adaptive multigrid solver. These developments allowed simulation of 2D and 3D viscous fingering and beginning analysis of systematic errors in the Todd-Longstaff mixing model. Trangenstein also developed modifications of the Colella corner transport upwind scheme for linear advection with a pre-specified divergence-free velocity field. Jeffrey Hittinger was very helpful in getting this research started. The analysis of the 3D stability conditions has not yet been completed, so this scheme has not been implemented yet.

Trangenstein made some useful modifications to the Deferred Execution Tool developed by Bill Allard at Duke, replacing some synchronous communications with asynchronous ones and reducing the number of messages needed in communications between patches. These changes reduced the communication costs by about 25%. Finally, he modified the interactive graphics for adaptive mesh refinement to operate with distributed processor arrays. This was somewhat tricky because 3D graphics makes greater use of callbacks for features such as image rotation and cut-planes.

Duke graduate student Wenjun Ying accompanied Trangenstein during the sabbatical. Ying worked on hybrid adaptive mesh refinement code for modeling electrical wave propagation in the heart. His approach is to maintain an array of grid cells within a given level of refinement, rather than an array of grid patches. This gives him somewhat greater flexibility to deal with heart geometry. Space-filling curves will be used to subdivide the grid cell array into sub-arrays on separate processors. During this visit, Ying programmed and debugged a 3D adaptive multigrid algorithm for the elliptic and parabolic equations in the cardiac bidomain model. Incorporation of reactions will follow.

*Continued*

### *Summary continued*

---

Through previous employment at LLNL and association with a former Ph.D. student (Richard Hornung) and a former collaborator (Xabier Garaizar) now working at LLNL, Trangenstein had good ties to LLNL before his sabbatical visit. There are several research projects that could help to continue such connections. In order to make adaptive mesh refinement more efficient for computations involving wells, it would be helpful to use cylindrical grids around the wells and Cartesian grids between the wells. Such an approach would require irregular communications between grid patches; several LLNL CASC personnel have experience in this approach.

Compressible multi-phase flow in porous media involves nonlinear interactions between mass-conservation and the velocity/pressure field. The solution of these nonlinear equations is expected to completely dominate the computational work in black-oil models and would need to be implemented efficiently in adaptive mesh refinement. Communication with CASC member Peter Brown could be very helpful in this work.

CASC post-doctoral fellow Miguel Dumett, soon to become assistant professor at the University of Southern California, is interested in developing a collaboration with John Trangenstein in work on modeling electrical wave propagation in the heart.

# Two-Grid Operators and Coarse Space Construction in Algebraic Multigrid Methods

**Ludmil Zikatanov**

&

**Rolf Ryham**

Pennsylvania State University

*LLNL Contact*

**Rob Falgout**

## Summary

---

The main ingredients in constructing an algebraic multigrid methods construction are building a hierarchy of coarse spaces (coarsening) and applying an approximate solver in each of these spaces (smoothing). The development of an efficient and robust algorithm applicable to wide class of problems requires careful numerical and theoretical investigation of the relationships between approximations in different norms and local-global behavior of linear operators (obtained by various discretization techniques). In collaboration with Rob Falgout and Panayot Vassilevski from Livermore's Center for Advanced Scientific Computing, we have worked on the convergence analysis for two-grid methods aiming to apply the theoretical results in the design of more efficient algebraic multigrid methods. We have derived a sharp bound on the convergence for a general type of two-level product iterative methods (in fact, we have an identity for the norm of the error transfer operator for these methods, not just the bound). Such a result not only justifies the use of various quality measures used in the coarsening phase, but in most cases, it also provides necessary and sufficient conditions for uniform (say, with respect to mesh parameters) convergence.

We have also studied (and plan to continue to study) several coarsening algorithms based on this estimate. The idea is to use the smoother to find first a subspace, which is complementary to a coarse space, and then construct the coarse space by using a minimization of appropriate energy functional. This approach allows the use of a prescribed sparsity pattern of the interpolation matrix and an adaptive choice of coarse spaces as well. The possible applications of the techniques we have studied are the algebraic multigrid methods and some of the numerical homogenization methods. We plan to test these numerical solution algorithms by applying them to discretized systems of partial differential equations, such as linear elasticity and Maxwell equations, and to further study the possibility of building smoothers and coarse spaces iteratively, i.e., not only vary and adopt the hierarchy of coarse spaces but also appropriately change the smoother during multigrid iterations.